

(12) **United States Patent**
Archer et al.

(10) **Patent No.:** **US 9,189,289 B2**
(45) **Date of Patent:** ***Nov. 17, 2015**

(54) **EXECUTING A COLLECTIVE OPERATION ALGORITHM IN A PARALLEL COMPUTER**

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION**,
Armonk, NY (US)

(72) Inventors: **Charles J. Archer**, Rochester, MN (US);
James E. Carey, Rochester, MN (US);
Philip J. Sanders, Rochester, MN (US);
Brian E. Smith, Knoxville, TN (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 339 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **13/766,164**

(22) Filed: **Feb. 13, 2013**

(65) **Prior Publication Data**

US 2014/0165076 A1 Jun. 12, 2014

Related U.S. Application Data

(63) Continuation of application No. 13/706,694, filed on Dec. 6, 2012.

(51) **Int. Cl.**
G06F 9/455 (2006.01)
G06F 9/46 (2006.01)
G06F 9/50 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 9/5066** (2013.01)

(58) **Field of Classification Search**

None

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,958,183 B2 *	6/2011	Arimilli et al.	709/201
2008/0263329 A1 *	10/2008	Archer et al.	712/213
2009/0043910 A1 *	2/2009	Barsness et al.	709/237
2010/0100703 A1 *	4/2010	Basu et al.	712/11
2010/0293269 A1 *	11/2010	Wilson et al.	709/224
2011/0191437 A1	8/2011	Chen et al.	
2011/0270986 A1	11/2011	Archer et al.	

* cited by examiner

Primary Examiner — Lewis A Bullock, Jr.

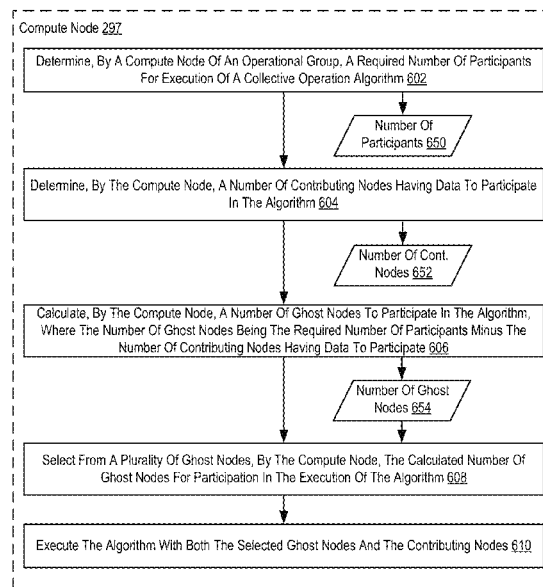
Assistant Examiner — Melissa Alfred

(74) *Attorney, Agent, or Firm* — Edward J. Lenart; Grant A. Johnson; Kennedy Lenart Spraggins LLP

(57) **ABSTRACT**

Executing a collective operation algorithm in a parallel computer includes a compute node of an operational group determining a required number of participants for execution of a collective operation algorithm and determining a number of contributing nodes having data to participate in the algorithm. Embodiments also include the compute node calculating a number of ghost nodes to participate in the algorithm. According to embodiments of the present invention, the number of ghost nodes is the required number of participants minus the number of contributing nodes having data to participate. Embodiments also include the compute node selecting from a plurality of ghost nodes, the calculated number of ghost nodes for participation in the execution of the algorithm and executing the algorithm with both the selected ghost nodes and the contributing nodes.

4 Claims, 10 Drawing Sheets



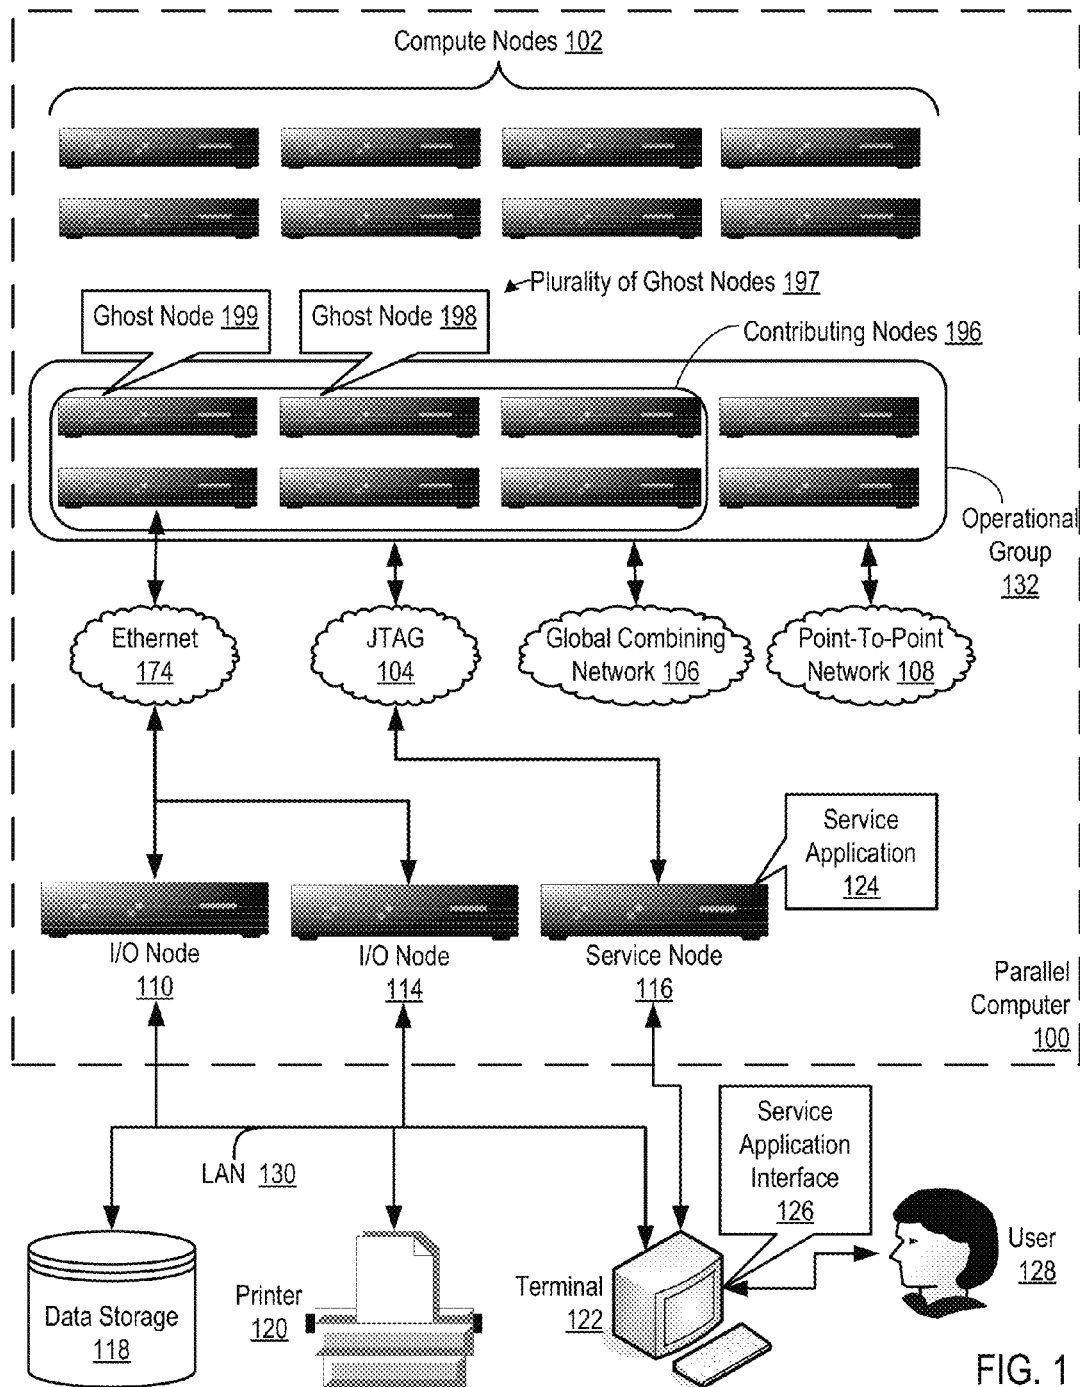



FIG. 1






FIG. 3A





FIG. 3B



Point-To-Point Network, Organized As A
'Torus' Or 'Mesh' 108

FIG. 4

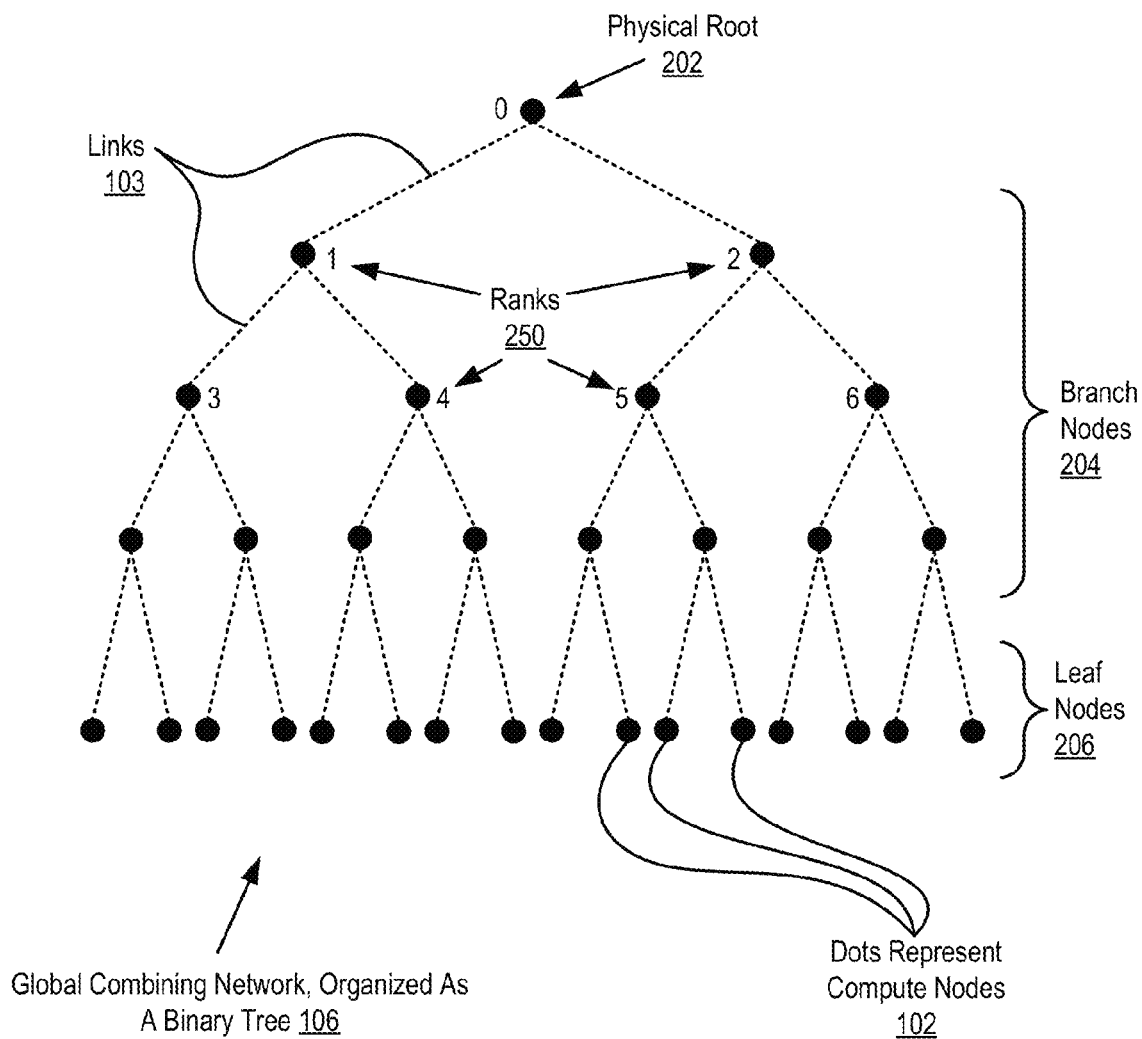


FIG. 5

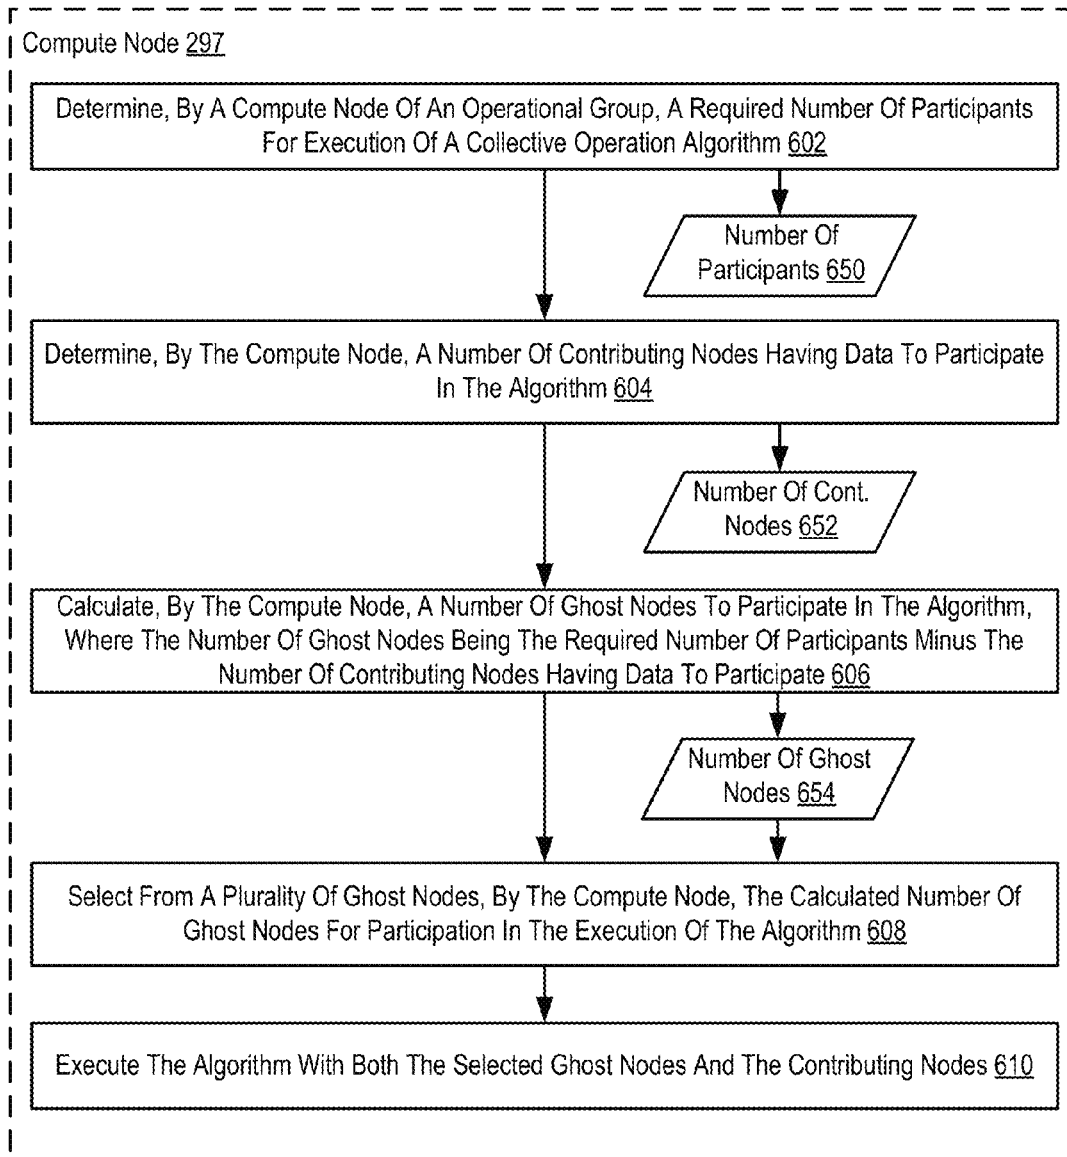


FIG. 6

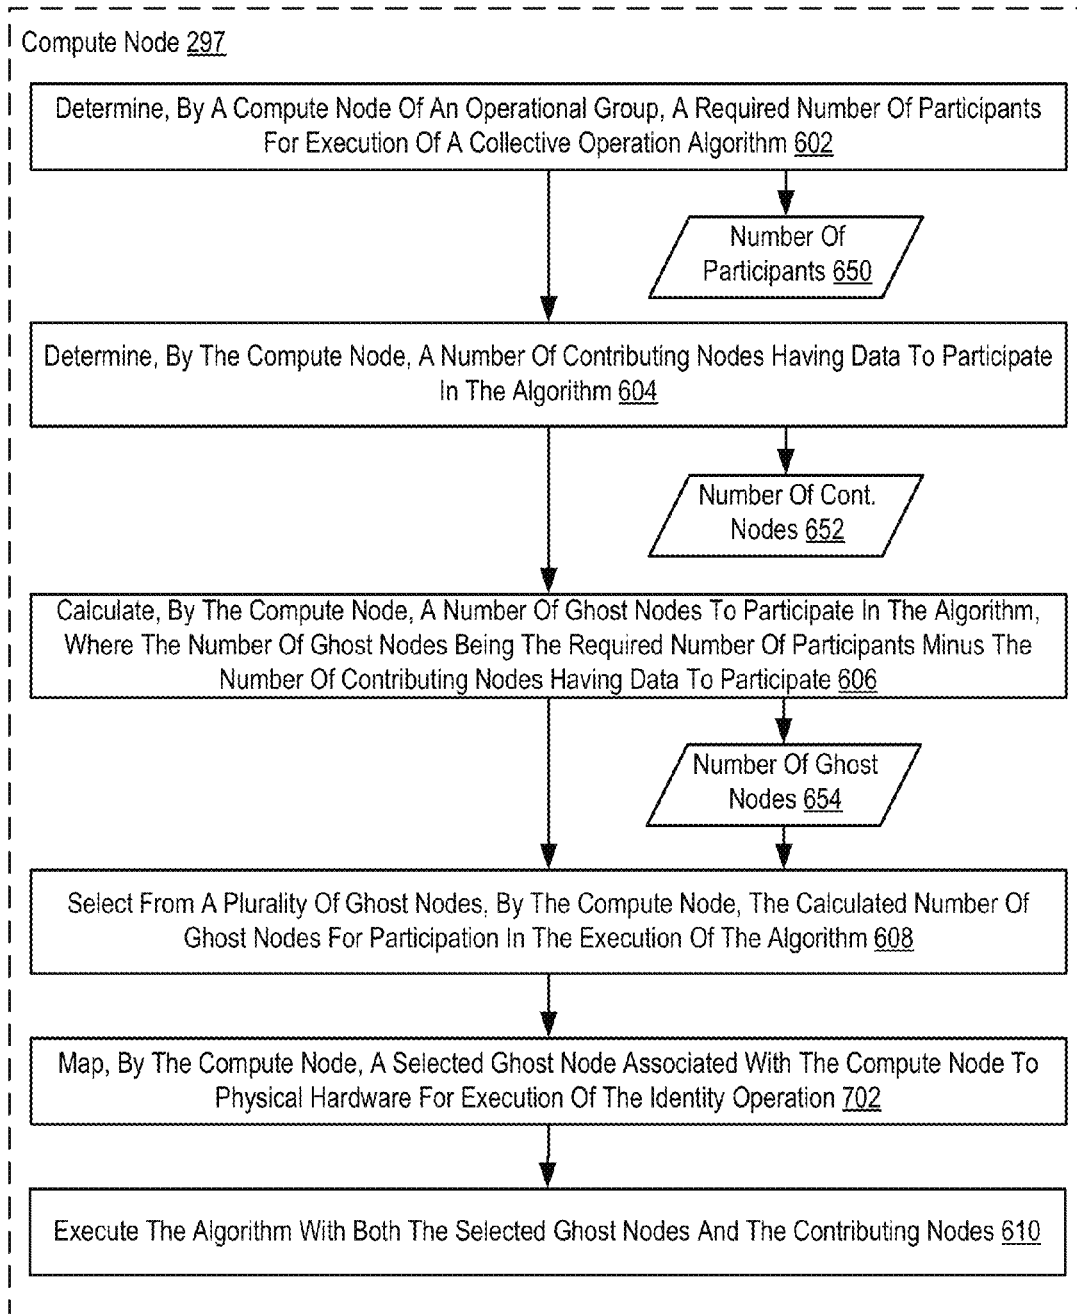


FIG. 7

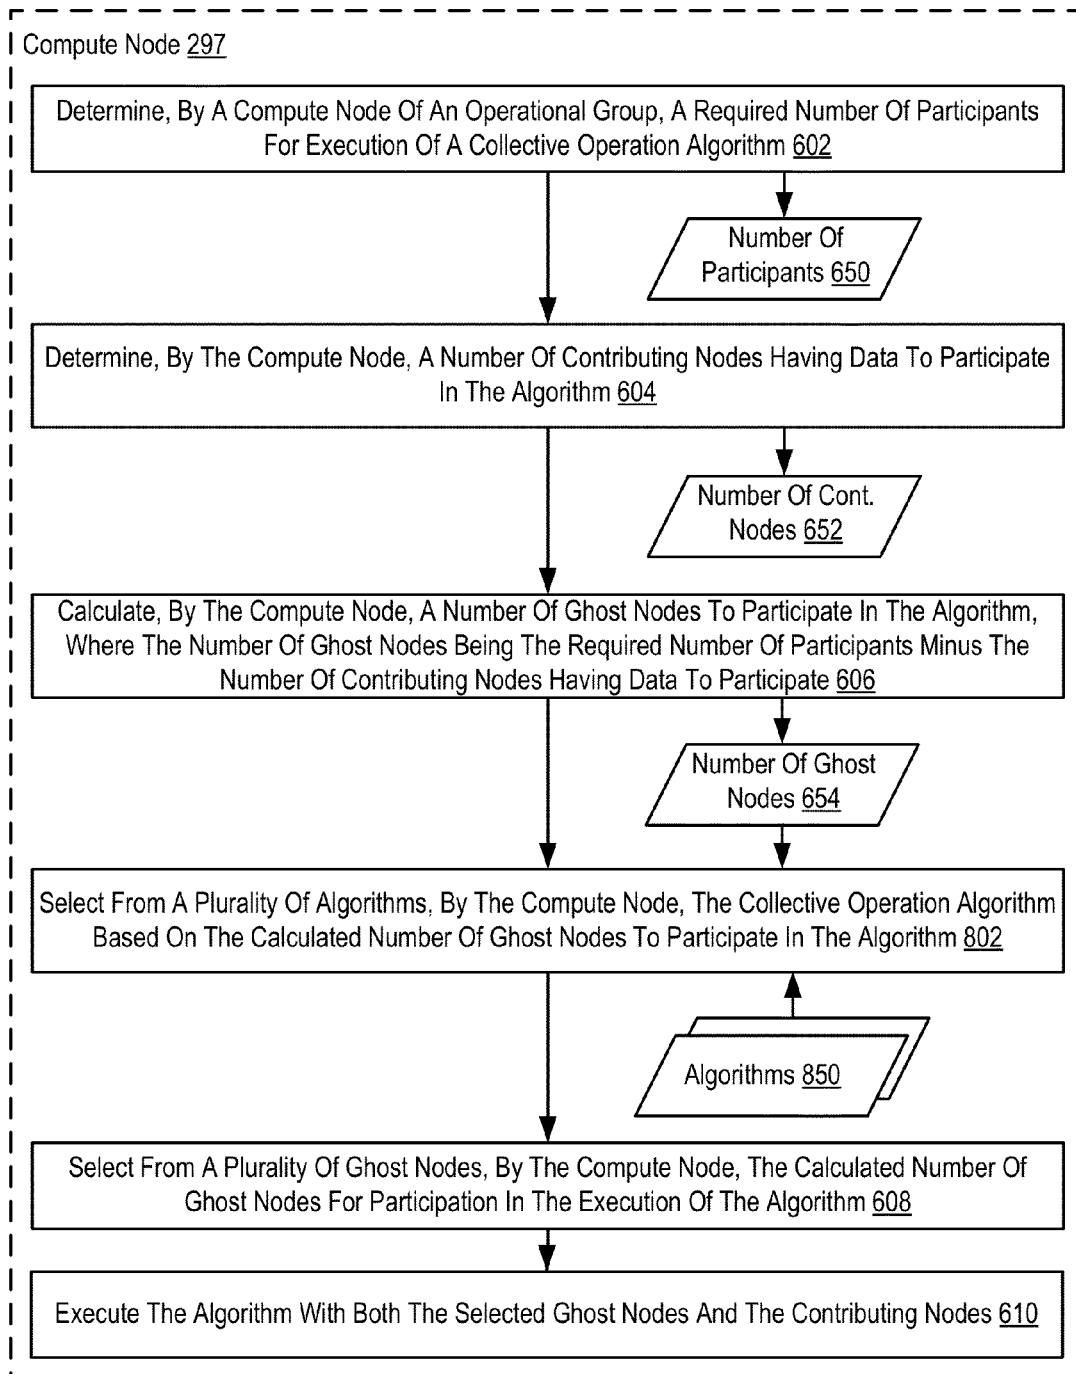


FIG. 8

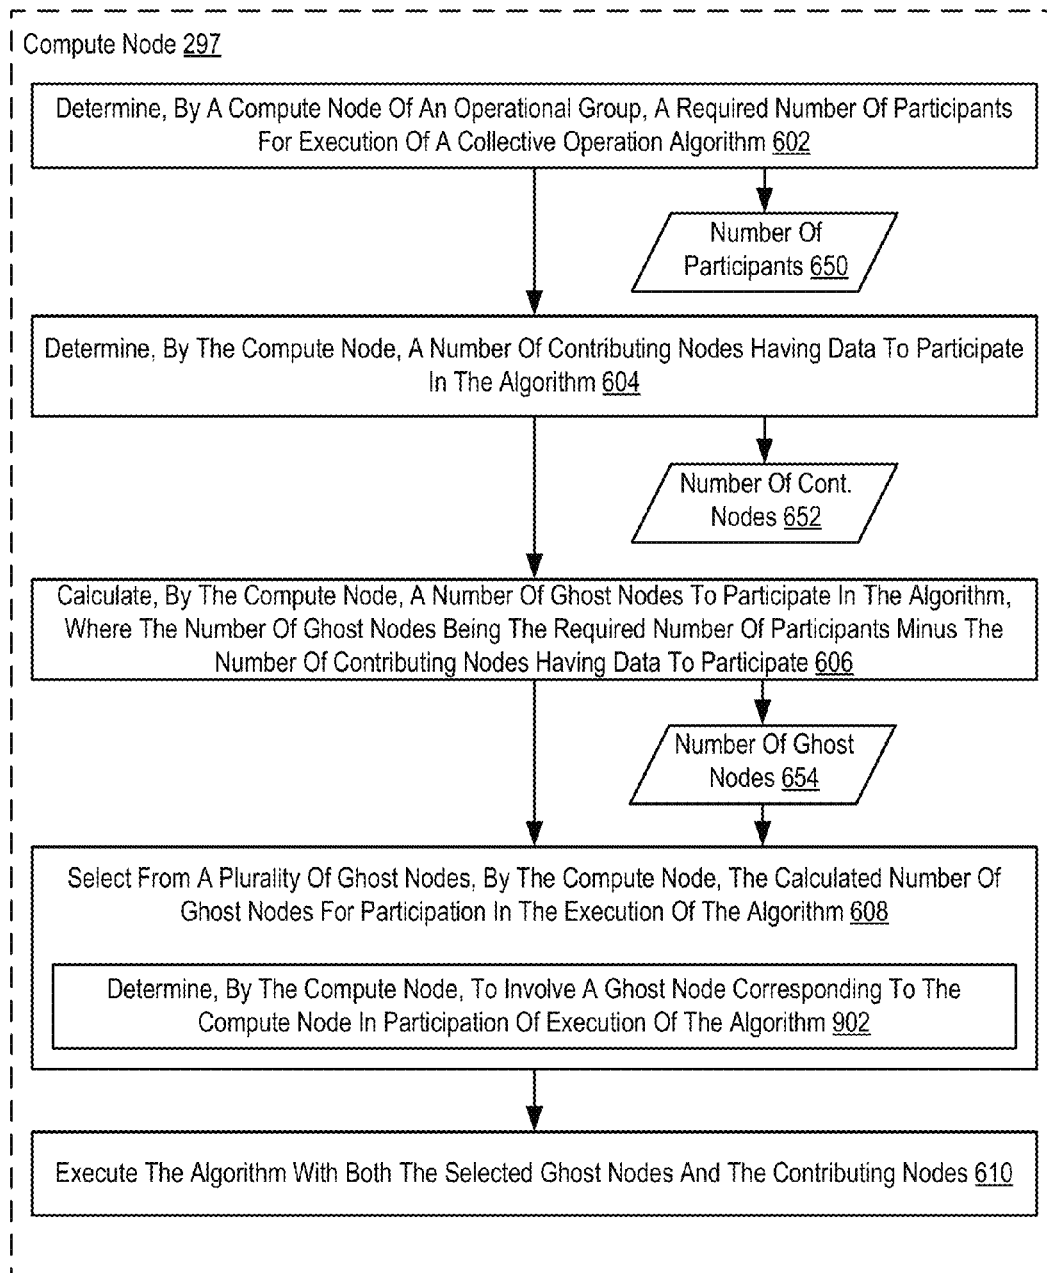


FIG. 9

1

EXECUTING A COLLECTIVE OPERATION ALGORITHM IN A PARALLEL COMPUTER

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation application of and claims priority from U.S. patent application Ser. No. 13/706,694, filed on Dec. 6, 2012.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The field of the invention is data processing, or, more specifically, methods, apparatus, and products for executing a collective operation algorithm in a parallel computer.

2. Description of Related Art

The development of the EDVAC computer system of 1948 is often cited as the beginning of the computer era. Since that time, computer systems have evolved into extremely complicated devices. Today's computers are much more sophisticated than early systems such as the EDVAC. Computer systems typically include a combination of hardware and software components, application programs, operating systems, processors, buses, memory, input/output devices, and so on. As advances in semiconductor processing and computer architecture push the performance of the computer higher and higher, more sophisticated computer software has evolved to take advantage of the higher performance of the hardware, resulting in computer systems today that are much more powerful than just a few years ago.

Parallel computing is an area of computer technology that has experienced advances. Parallel computing is the simultaneous execution of the same task (split up and specially adapted) on multiple processors in order to obtain results faster. Parallel computing is based on the fact that the process of solving a problem usually can be divided into smaller tasks, which may be carried out simultaneously with some coordination.

Parallel computers execute parallel algorithms. A parallel algorithm can be split up to be executed a piece at a time on many different processing devices, and then put back together again at the end to get a data processing result. Some algorithms are easy to divide up into pieces. Splitting up the job of checking all of the numbers from one to a hundred thousand to see which are primes could be done, for example, by assigning a subset of the numbers to each available processor, and then putting the list of positive results back together. In this specification, the multiple processing devices that execute the individual pieces of a parallel program are referred to as 'compute nodes.' A parallel computer is composed of compute nodes and other processing nodes as well, including, for example, input/output ('I/O') nodes, and service nodes.

Parallel algorithms are valuable because it is faster to perform some kinds of large computing tasks via a parallel algorithm than it is via a serial (non-parallel) algorithm, because of the way modern processors work. It is far more difficult to construct a computer with a single fast processor than one with many slow processors with the same throughput. There are also certain theoretical limits to the potential speed of serial processors. On the other hand, every parallel algorithm has a serial part and so parallel algorithms have a saturation point. After that point adding more processors does not yield any more throughput but only increases the overhead and cost.

2

Parallel algorithms are designed also to optimize one more resource the data communications requirements among the nodes of a parallel computer. There are two ways parallel processors communicate, shared memory or message passing. Shared memory processing needs additional locking for the data and imposes the overhead of additional processor and bus cycles and also serializes some portion of the algorithm.

Message passing processing uses high-speed data communications networks and message buffers, but this communication adds transfer overhead on the data communications networks as well as additional memory need for message buffers and latency in the data communications among nodes. Designs of parallel computers use specially designed data communications links so that the communication overhead will be small but it is the parallel algorithm that decides the volume of the traffic.

Many data communications network architectures are used for message passing among nodes in parallel computers. Compute nodes may be organized in a network as a 'torus' or 'mesh,' for example. Also, compute nodes may be organized in a network as a tree. A torus network connects the nodes in a three-dimensional mesh with wrap around links. Every node is connected to its six neighbors through this torus network, and each node is addressed by its x,y,z coordinate in the mesh. In such a manner, a torus network lends itself to point to point operations. In a tree network, the nodes typically are connected into a binary tree: each node has a parent, and two children (although some nodes may only have zero children or one child, depending on the hardware configuration). Although a tree network typically is inefficient in point to point communication, a tree network does provide high bandwidth and low latency for certain collective operations, message passing operations where all compute nodes participate simultaneously, such as, for example, an allgather operation. In computers that use a torus and a tree network, the two networks typically are implemented independently of one another, with separate routing circuits, separate physical links, and separate message buffers.

SUMMARY OF THE INVENTION

Methods, apparatuses, and computer program products for executing a collective operation algorithm in a parallel computer are disclosed in this specification. The parallel computer includes a plurality of compute nodes coupled for data communications by one or more data communications networks. The compute nodes are organized in an operational group for collective operations. Executing a collective operation algorithm in a parallel computer in accordance with embodiments of the present invention includes a compute node of an operational group determining a required number of participants for execution of a collective operation algorithm and determining a number of contributing nodes having data to participate in the algorithm. Embodiments also include the compute node calculating a number of ghost nodes to participate in the algorithm. According to embodiments of the present invention, the calculated number of ghost nodes is the required number of participants minus the number of contributing nodes having data to participate. Embodiments also include the compute node selecting from a plurality of ghost nodes, the calculated number of ghost nodes for participation in the execution of the algorithm and executing the algorithm with both the selected ghost nodes and the contributing nodes.

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular descriptions of exemplary embodiments of the

invention as illustrated in the accompanying drawings wherein like reference numbers generally represent like parts of exemplary embodiments of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an exemplary system for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 2 sets forth a block diagram of an example compute node useful in a parallel computer capable of initiating a collective operation according to embodiments of the present invention.

FIG. 3A sets forth a block diagram of an example Point-To-Point Adapter useful in systems for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 3B sets forth a block diagram of an example Global Combining Network Adapter useful in systems for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 4 sets forth a line drawing illustrating an example data communications network optimized for point-to-point operations useful in systems capable of executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 5 sets forth a line drawing illustrating an example global combining network useful in systems capable of executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 6 sets forth a flow chart illustrating an example method for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 7 sets forth a flow chart illustrating another example method for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 8 sets forth a flow chart illustrating another example method for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

FIG. 9 sets forth a flow chart illustrating another example method for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

Exemplary methods, apparatuses, and computer program products for executing a collective operation algorithm in a parallel computer in accordance with the present invention are described with reference to the accompanying drawings, beginning with FIG. 1. FIG. 1 illustrates an exemplary system for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The system of FIG. 1 includes a parallel computer (100), non-volatile memory for the computer in the form of a data storage device (118), an output device for the computer in the form of a printer (120), and an input/output device for the computer in the form of a computer terminal (122).

The parallel computer (100) in the example of FIG. 1 includes a plurality of compute nodes (102). The compute nodes (102) are coupled for data communications by several independent data communications networks including a high speed Ethernet network (174), a Joint Test Action Group

(‘JTAG’) network (104), a global combining network (106) which is optimized for collective operations using a binary tree network topology, and a point-to-point network (108), which is optimized for point-to-point operations using a torus network topology. The global combining network (106) is a data communications network that includes data communications links connected to the compute nodes (102) so as to organize the compute nodes (102) as a binary tree. Each data communications network is implemented with data communications links among the compute nodes (102). The data communications links provide data communications for parallel operations among the compute nodes (102) of the parallel computer (100).

The compute nodes (102) of the parallel computer (100) are organized into at least one operational group (132) of compute nodes for collective parallel operations on the parallel computer (100). Each operational group (132) of compute nodes is the set of compute nodes upon which a collective parallel operation executes. Each compute node in the operational group (132) is assigned a unique rank that identifies the particular compute node in the operational group (132). Collective operations are implemented with data communications among the compute nodes of an operational group. Collective operations are those functions that involve all the compute nodes of an operational group (132). A collective operation is an operation, a message-passing computer program instruction that is executed simultaneously, that is, at approximately the same time, by all the compute nodes in an operational group (132) of compute nodes. Such an operational group (132) may include all the compute nodes (102) in a parallel computer (100) or a subset all the compute nodes (102). Collective operations are often built around point-to-point operations. A collective operation requires that all processes on all compute nodes within an operational group (132) call the same collective operation with matching arguments. A ‘broadcast’ is an example of a collective operation for moving data among compute nodes of an operational group. A ‘reduce’ operation is an example of a collective operation that executes arithmetic or logical functions on data distributed among the compute nodes of an operational group (132). An operational group (132) may be implemented as, for example, an MPI ‘communicator.’

‘MPI’ refers to ‘Message Passing Interface,’ a prior art parallel communications library, a module of computer program instructions for data communications on parallel computers. Examples of prior-art parallel communications libraries that may be improved for use in systems configured according to embodiments of the present invention include MPI and the ‘Parallel Virtual Machine’ (‘PVM’) library. PVM was developed by the University of Tennessee, The Oak Ridge National Laboratory and Emory University. MPI is promulgated by the MPI Forum, an open group with representatives from many organizations that define and maintain the MPI standard. MPI at the time of this writing is a de facto standard for communication among compute nodes running a parallel program on a distributed memory parallel computer. This specification sometimes uses MPI terminology for ease of explanation, although the use of MPI as such is not a requirement or limitation of the present invention.

Some collective operations have a single originating or receiving process running on a particular compute node in an operational group (132). For example, in a ‘broadcast’ collective operation, the process on the compute node that distributes the data to all the other compute nodes is an originating process. In a ‘gather’ operation, for example, the process on the compute node that received all the data from the other

compute nodes is a receiving process. The compute node on which such an originating or receiving process runs is referred to as a logical root.

Most collective operations are variations or combinations of four basic operations: broadcast, gather, scatter, and reduce. The interfaces for these collective operations are defined in the MPI standards promulgated by the MPI Forum. Algorithms for executing collective operations, however, are not defined in the MPI standards. In a broadcast operation, all processes specify the same root process, whose buffer contents will be sent. Processes other than the root specify receive buffers. After the operation, all buffers contain the message from the root process.

A scatter operation, like the broadcast operation, is also a one-to-many collective operation. In a scatter operation, the logical root divides data on the root into segments and distributes a different segment to each compute node in the operational group (132). In scatter operation, all processes typically specify the same receive count. The send arguments are only significant to the root process, whose buffer actually contains sendcount*N elements of a given datatype, where N is the number of processes in the given group of compute nodes. The send buffer is divided and dispersed to all processes (including the process on the logical root). Each compute node is assigned a sequential identifier termed a 'rank.' After the operation, the root has sent sendcount data elements to each process in increasing rank order. Rank 0 receives the first sendcount data elements from the send buffer. Rank 1 receives the second sendcount data elements from the send buffer, and so on.

A gather operation is a many-to-one collective operation that is a complete reverse of the description of the scatter operation. That is, a gather is a many-to-one collective operation in which elements of a datatype are gathered from the ranked compute nodes into a receive buffer in a root node.

A reduction operation is also a many-to-one collective operation that includes an arithmetic or logical function performed on two data elements. All processes specify the same 'count' and the same arithmetic or logical function. After the reduction, all processes have sent count data elements from compute node send buffers to the root process. In a reduction operation, data elements from corresponding send buffer locations are combined pair-wise by arithmetic or logical operations to yield a single corresponding element in the root process' receive buffer. Application specific reduction operations can be defined at runtime. Parallel communications libraries may support predefined operations. MPI, for example, provides the following predefined reduction operations:

MPI_MAX	maximum
MPI_MIN	minimum
MPI_SUM	sum
MPI_PROD	product
MPI_LAND	logical and
MPI_BAND	bitwise and
MPI_LOR	logical or
MPI BOR	bitwise or
MPI_LXOR	logical exclusive or
MPI_BXOR	bitwise exclusive or

In addition to compute nodes, the parallel computer (100) includes input/output ('I/O') nodes (110, 114) coupled to compute nodes (102) through the global combining network (106). The compute nodes (102) in the parallel computer (100) may be partitioned into processing sets such that each compute node in a processing set is connected for data com-

munications to the same I/O node. Each processing set, therefore, is composed of one I/O node and a subset of compute nodes (102). The ratio between the number of compute nodes to the number of I/O nodes in the entire system typically depends on the hardware configuration for the parallel computer (102). For example, in some configurations, each processing set may be composed of eight compute nodes and one I/O node. In some other configurations, each processing set may be composed of sixty-four compute nodes and one I/O node. Such examples are for explanation only, however, and not for limitation. Each I/O node provides I/O services between compute nodes (102) of its processing set and a set of I/O devices. In the example of FIG. 1, the I/O nodes (110, 114) are connected for data communications I/O devices (118, 120, 122) through local area network ('LAN') (130) implemented using high-speed Ethernet.

The parallel computer (100) of FIG. 1 also includes a service node (116) coupled to the compute nodes through one of the networks (104). Service node (116) provides services common to pluralities of compute nodes, administering the configuration of compute nodes, loading programs into the compute nodes, starting program execution on the compute nodes, retrieving results of program operations on the compute nodes, and so on. Service node (116) runs a service application (124) and communicates with users (128) through a service application interface (126) that runs on computer terminal (122).

The parallel computer (100) of FIG. 1 operates generally for executing a collective operation algorithm in a parallel computer in accordance with embodiments of the present invention. The compute nodes (102) of the example parallel computer (100) are organized for in the operational group (132) for executing collective operation algorithms.

A collective operation algorithm is designed for multiple compute nodes to execute pieces of a program in parallel. Some collective operation algorithms are very efficient but may require a particular number of participants. Examples of such algorithms include algorithms that require a number of participants that are a power of two, a multiple of two, an odd number of participants, a prime number of participants and so on.

In the example parallel computer (100) of FIG. 1, the compute nodes operate in accordance with embodiments of the present invention to execute a collective operation algorithm by determining a required number of participants for execution of a collective operation algorithm and determining a number of contributing nodes (196) having data to participate in the algorithm. Embodiments also include the compute nodes calculating a number of ghost nodes (199, 198) to participate in the algorithm.

A ghost node is a virtual node configured to participate in the collective operation algorithm, such as with an identity operation. That is, the ghost node is a construct mapped to an actual node that may contribute an identity operation to the execution of the collective operation algorithm. The identity operation is a function that always returns the same value that was used as its argument. An example identity equation is $f(x)=x$. The identity operation for addition is plus zero and the identity operation for multiplication is multiplication by one.

According to embodiments of the present invention, the number of ghost nodes is the required number of participants minus the number of contributing nodes having data to participate. Embodiments also include the compute nodes selecting from a plurality of ghost nodes, the calculated number of ghost nodes for participation in the execution of the algorithm and executing the algorithm with both the selected ghost nodes and the contributing nodes. That is, according to

embodiments of the present invention, a compute node may execute a collective operation algorithm that requires a particular number of participants with an arbitrary number of compute nodes and a calculated number of ghost nodes.

Executing a collective operation algorithm according to embodiments of the present invention is generally implemented on a parallel computer that includes a plurality of compute nodes organized for collective operations through at least one data communications network. In fact, such computers may include thousands of such compute nodes. Each compute node is in turn itself a kind of computer composed of one or more computer processing cores, its own computer memory, and its own input/output adapters. For further explanation, therefore, FIG. 2 sets forth a block diagram of an example compute node (297) useful in a parallel computer capable of initiating a collective operation according to embodiments of the present invention. The compute node (297) of FIG. 2 includes a plurality of processing cores (165) as well as RAM (156). The processing cores (165) of FIG. 2 may be configured on one or more integrated circuit dies. Processing cores (165) are connected to RAM (156) through a high-speed memory bus (155) and through a bus adapter (194) and an extension bus (168) to other components of the compute node.

Stored in RAM (156) is an application program (226), a module of computer program instructions that carries out parallel, user-level data processing using parallel algorithms, such as collective operation algorithm (228). In the example of FIG. 2, the application (226) may implement a participant in an operational group, such as a rank in an MPI-style communicator. Execution of the application program (226) causes the example compute node (297) of FIG. 2 to execute the collective operation algorithm (228) in accordance with embodiments of the present invention. The compute node (297) may carry out such collective operation algorithm execution by: determining a required number of participants for execution of a collective operation algorithm (228) and determining a number of contributing nodes having data to participate in the algorithm (228). Embodiments also include the compute node calculating a number of ghost nodes to participate in the algorithm. A ghost node is a virtual node configured to participate in the collective operation algorithm, such as with an identity operation. In the example of FIG. 2, the compute node has an associated ghost node (298). That is, the ghost node is a construct mapped to an actual node that may contribute an identity operation (299) to the execution of the collective operation algorithm. The identity operation (299) is a function that always returns the same value that was used as its argument. According to embodiments of the present invention, the calculated number of ghost nodes is the required number of participants minus the number of contributing nodes having data to participate. Embodiments also include the compute node selecting from a plurality of ghost nodes, the calculated number of ghost nodes for participation in the execution of the algorithm (228) and executing the algorithm (228) with both the selected ghost nodes and the contributing nodes.

Also stored RAM (156) is a parallel communications library (161), a library of computer program instructions that carry out parallel communications among compute nodes, including point-to-point operations as well as collective operations. A library of parallel communications routines may be developed from scratch for use in systems according to embodiments of the present invention, using a traditional programming language such as the C programming language, and using traditional programming methods to write parallel communications routines that send and receive data among

nodes on two independent data communications networks. Alternatively, existing prior art libraries may be improved to operate according to embodiments of the present invention. Examples of prior-art parallel communications libraries include the 'Message Passing Interface' ('MPI') library and the 'Parallel Virtual Machine' ('PVM') library.

Also stored in RAM (156) is an operating system (162), a module of computer program instructions and routines for an application program's access to other resources of the compute node. It is typical for an application program and parallel communications library in a compute node of a parallel computer to run a single thread of execution with no user login and no security issues because the thread is entitled to complete access to all resources of the node. The quantity and complexity of tasks to be performed by an operating system on a compute node in a parallel computer therefore are smaller and less complex than those of an operating system on a serial computer with many threads running simultaneously. In addition, there is no video I/O on the compute node (297) of FIG. 2, another factor that decreases the demands on the operating system. The operating system (162) may therefore be quite lightweight by comparison with operating systems of general purpose computers, a pared down version as it were, or an operating system developed specifically for operations on a particular parallel computer. Operating systems that may usefully be improved, simplified, for use in a compute node include UNIX™, Linux™, Windows XP™, AIX™, IBM's i5/OS™, and others as will occur to those of skill in the art.

The example compute node (297) of FIG. 2 includes several communications adapters (172, 176, 180, 188) for implementing data communications with other nodes of a parallel computer. Such data communications may be carried out serially through RS-232 connections, through external buses such as USB, through data communications networks such as IP networks, and in other ways as will occur to those of skill in the art. Communications adapters implement the hardware level of data communications through which one computer sends data communications to another computer, directly or through a network. Examples of communications adapters useful in apparatus useful for executing a collective operation algorithm in a parallel computer include modems for wired communications, Ethernet (IEEE 802.3) adapters for wired network communications, and 802.11b adapters for wireless network communications.

The data communications adapters in the example of FIG. 2 include a Gigabit Ethernet adapter (172) that couples example compute node (297) for data communications to a Gigabit Ethernet (174). Gigabit Ethernet is a network transmission standard, defined in the IEEE 802.3 standard, that provides a data rate of 1 billion bits per second (one gigabit). Gigabit Ethernet is a variant of Ethernet that operates over multimode fiber optic cable, single mode fiber optic cable, or unshielded twisted pair.

The data communications adapters in the example of FIG. 2 include a JTAG Slave circuit (176) that couples example compute node (297) for data communications to a JTAG Master circuit (178). JTAG is the usual name used for the IEEE 1149.1 standard entitled Standard Test Access Port and Boundary-Scan Architecture for test access ports used for testing printed circuit boards using boundary scan. JTAG is so widely adapted that, at this time, boundary scan is more or less synonymous with JTAG. JTAG is used not only for printed circuit boards, but also for conducting boundary scans of integrated circuits, and is also useful as a mechanism for debugging embedded systems, providing a convenient alternative access point into the system. The example compute node of FIG. 2 may be all three of these: It typically includes

one or more integrated circuits installed on a printed circuit board and may be implemented as an embedded system having its own processing core, its own memory, and its own I/O capability. JTAG boundary scans through JTAG Slave (176) may efficiently configure processing core registers and memory in compute node (297) for use in dynamically reassigning a connected node to a block of compute nodes useful in systems for executing a collective operation algorithm in a parallel computer to embodiments of the present invention.

The data communications adapters in the example of FIG. 2 include a Point-To-Point Network Adapter (180) that couples example compute node (297) for data communications to a network (108) that is optimal for point-to-point message passing operations such as, for example, a network configured as a three-dimensional torus or mesh. The Point-To-Point Adapter (180) provides data communications in six directions on three communications axes, x, y, and z, through six bidirectional links: +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186).

The data communications adapters in the example of FIG. 2 include a Global Combining Network Adapter (188) that couples example compute node (297) for data communications to a global combining network (106) that is optimal for collective message passing operations such as, for example, a network configured as a binary tree. The Global Combining Network Adapter (188) provides data communications through three bidirectional links for each global combining network (106) that the Global Combining Network Adapter (188) supports. In the example of FIG. 2, the Global Combining Network Adapter (188) provides data communications through three bidirectional links for global combining network (106): two to children nodes (190) and one to a parent node (192).

The example compute node (297) includes multiple arithmetic logic units ('ALUs'). Each processing core (165) includes an ALU (166), and a separate ALU (170) is dedicated to the exclusive use of the Global Combining Network Adapter (188) for use in performing the arithmetic and logical functions of reduction operations, including an all reduce operation. Computer program instructions of a reduction routine in a parallel communications library (161) may latch an instruction for an arithmetic or logical function into an instruction register (169). When the arithmetic or logical function of a reduction operation is a 'sum' or a 'logical OR,' for example, the collective operations adapter (188) may execute the arithmetic or logical operation by use of the ALU (166) in the processing core (165) or, typically much faster, by use of the dedicated ALU (170) using data provided by the nodes (190, 192) on the global combining network (106) and data provided by processing cores (165) on the compute node (297).

Often when performing arithmetic operations in the global combining network adapter (188), however, the global combining network adapter (188) only serves to combine data received from the children nodes (190) and pass the result up the network (106) to the parent node (192). Similarly, the global combining network adapter (188) may only serve to transmit data received from the parent node (192) and pass the data down the network (106) to the children nodes (190). That is, none of the processing cores (165) on the compute node (297) contribute data that alters the output of ALU (170), which is then passed up or down the global combining network (106). Because the ALU (170) typically does not output any data onto the network (106) until the ALU (170) receives input from one of the processing cores (165), a processing core (165) may inject the identity element into the dedicated ALU (170) for the particular arithmetic operation being per-

form in the ALU (170) in order to prevent alteration of the output of the ALU (170). Injecting the identity element into the ALU, however, often consumes numerous processing cycles. To further enhance performance in such cases, the example compute node (297) includes dedicated hardware (171) for injecting identity elements into the ALU (170) to reduce the amount of processing core resources required to prevent alteration of the ALU output. The dedicated hardware (171) injects an identity element that corresponds to the particular arithmetic operation performed by the ALU. For example, when the global combining network adapter (188) performs a bitwise OR on the data received from the children nodes (190), dedicated hardware (171) may inject zeros into the ALU (170) to improve performance throughout the global combining network (106).

For further explanation, FIG. 3A sets forth a block diagram of an example Point-To-Point Adapter (180) useful in systems for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The Point-To-Point Adapter (180) is designed for use in a data communications network optimized for point-to-point operations, a network that organizes compute nodes in a three-dimensional torus or mesh. The Point-To-Point Adapter (180) in the example of FIG. 3A provides data communication along an x-axis through four unidirectional data communications links, to and from the next node in the -x direction (182) and to and from the next node in the +x direction (181). The Point-To-Point Adapter (180) of FIG. 3A also provides data communication along a y-axis through four unidirectional data communications links, to and from the next node in the -y direction (184) and to and from the next node in the +y direction (183). The Point-To-Point Adapter (180) of FIG. 3A also provides data communication along a z-axis through four unidirectional data communications links, to and from the next node in the -z direction (186) and to and from the next node in the +z direction (185). For further explanation, FIG. 3B sets forth a block diagram of an example Global Combining Network Adapter (188) useful in systems for executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The Global Combining Network Adapter (188) is designed for use in a network optimized for collective operations, a network that organizes compute nodes of a parallel computer in a binary tree. The Global Combining Network Adapter (188) in the example of FIG. 3B provides data communication to and from children nodes of a global combining network through four unidirectional data communications links (190), and also provides data communication to and from a parent node of the global combining network through two unidirectional data communications links (192).

For further explanation, FIG. 4 sets forth a line drawing illustrating an example data communications network (108) optimized for point-to-point operations useful in systems capable of executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. In the example of FIG. 4, dots represent compute nodes (102) of a parallel computer, and the dotted lines between the dots represent data communications links (103) between compute nodes. The data communications links are implemented with point-to-point data communications adapters similar to the one illustrated for example in FIG. 3A, with data communications links on three axes, x, y, and z, and to and fro in six directions +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186). The links and compute nodes are organized by this data communications network optimized for point-to-point operations into a three dimensional mesh (105). The mesh (105) has wrap-around links on each

11

axis that connect the outermost compute nodes in the mesh (105) on opposite sides of the mesh (105). These wrap-around links form a torus (107). Each compute node in the torus has a location in the torus that is uniquely specified by a set of x, y, z coordinates. Readers will note that the wrap-around links in the y and z directions have been omitted for clarity, but are configured in a similar manner to the wrap-around link illustrated in the x direction. For clarity of explanation, the data communications network of FIG. 4 is illustrated with only 27 compute nodes, but readers will recognize that a data communications network optimized for point-to-point operations for use in executing a collective operation algorithm in a parallel computer in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes. For ease of explanation, the data communications network of FIG. 4 is illustrated with only three dimensions, but readers will recognize that a data communications network optimized for point-to-point operations for use in executing a collective operation algorithm in a parallel computer in accordance with embodiments of the present invention may in fact be implemented in two dimensions, four dimensions, five dimensions, and so on. Several supercomputers now use five dimensional mesh or torus networks, including, for example, IBM's Blue Gene Q™.

For further explanation, FIG. 5 sets forth a line drawing illustrating an example global combining network (106) useful in systems capable of executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The example data communications network of FIG. 5 includes data communications links (103) connected to the compute nodes so as to organize the compute nodes as a tree. In the example of FIG. 5, dots represent compute nodes (102) of a parallel computer, and the dotted lines (103) between the dots represent data communications links between compute nodes. The data communications links are implemented with global combining network adapters similar to the one illustrated for example in FIG. 3B, with each node typically providing data communications to and from two children nodes and data communications to and from a parent node, with some exceptions. Nodes in the global combining network (106) may be characterized as a physical root node (202), branch nodes (204), and leaf nodes (206). The physical root (202) has two children but no parent and is so called because the physical root node (202) is the node physically configured at the top of the binary tree. The leaf nodes (206) each has a parent, but leaf nodes have no children. The branch nodes (204) each has both a parent and two children. The links and compute nodes are thereby organized by this data communications network optimized for collective operations into a binary tree (106). For clarity of explanation, the data communications network of FIG. 5 is illustrated with only 31 compute nodes, but readers will recognize that a global combining network (106) optimized for collective operations for use in executing a collective operation algorithm in a parallel computer in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes.

In the example of FIG. 5, each node in the tree is assigned a unit identifier referred to as a 'rank' (250). The rank actually identifies a task or process that is executing a parallel operation according to embodiments of the present invention. Using the rank to identify a node assumes that only one such task is executing on each node. To the extent that more than one participating task executes on a single node, the rank identifies the task as such rather than the node. A rank uniquely identifies a task's location in the tree network for use

12

in both point-to-point and collective operations in the tree network. The ranks in this example are assigned as integers beginning with 0 assigned to the root tasks or root node (202), 1 assigned to the first node in the second layer of the tree, 2 assigned to the second node in the second layer of the tree, 3 assigned to the first node in the third layer of the tree, 4 assigned to the second node in the third layer of the tree, and so on. For ease of illustration, only the ranks of the first three layers of the tree are shown here, but all compute nodes in the tree network are assigned a unique rank.

For further explanation, FIG. 6 sets forth a flow chart illustrating an example method executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. For ease of explanation, the description of FIG. 6 makes reference to example components of FIG. 1 and FIG. 2.

The method of FIG. 6 includes determining (602) a required number (650) of participants for execution of a collective operation algorithm (228). A collective operation algorithm is designed for multiple compute nodes to execute pieces of a program in parallel. Some collective operation algorithms are very efficient but may require a particular number of participants. Examples of such algorithms include algorithms that require a number of participants that are a power of two, a multiple of two, an odd number of participants, a prime number of participants and so on. Determining (602) a required number (650) of participants for execution of a collective operation algorithm (228) may be carried out by examining a particular algorithm for an indication of the required number of participants.

The method of FIG. 6 also includes determining (604) a number (652) of contributing nodes (196) having data to participate in the algorithm (228). Determining (604) a number (652) of contributing nodes (196) having data to participate in the algorithm (228) may be carried out by receiving user input indicating particular compute nodes as 'contributing nodes.' For example, a user may indicate a particular set of nodes to participate in a collective operation algorithm.

The method of FIG. 6 includes calculating (606) a number (656) of ghost nodes to participate in the algorithm (228). A ghost node is a virtual node configured to participate in the collective operation algorithm, such as with an identity operation. That is, the ghost node is a construct mapped to an actual node that may contribute an identity operation to the execution of the collective operation algorithm. The identity operation is a function that always returns the same value that was used as its argument. An example identity equation is $f(x)=x$. The identity operation for addition is plus zero and the identity operation for multiplication is multiplication by one. According to embodiments of the present invention, the number of ghost nodes is the required number of participants minus the number of contributing nodes having data to participate. Calculating (606) a number (656) of ghost nodes to participate in the algorithm (228) may be carried out by subtracting the number (652) of contributing nodes from the required number (650) of participants. That is, the number of ghost nodes is the required number (650) of participants minus the number (652) of contributing nodes having data to participate.

The method of FIG. 6 also includes selecting (608) from a plurality of ghost nodes (197), the calculated number (654) of ghost nodes (198, 199) for participation in the execution of the algorithm (228). Any number of ghost nodes may be created and on any number of physical nodes. For example, a particular physical node may have one or more corresponding virtual or ghost nodes. Ghost nodes may be created at any time. For example, at job startup, each compute node may

13

create a ghost node of itself so there is a pool of available ghost nodes. Selecting (608) from a plurality of ghost nodes (197), the calculated number (654) of ghost nodes (198, 199) for participation in the execution of the algorithm (228) may be carried out by a compute node involving its associated ghost node in execution of the algorithm.

The method of FIG. 6 includes executing (610) the algorithm (228) with both the selected ghost nodes (198, 199) and the contributing nodes (196). Executing (610) the algorithm (228) with both the selected ghost nodes (198, 199) and the contributing nodes (196) may be carried out by each of the ghost nodes participating with an identity operation and the contributing nodes performing operations associated with the collective operation algorithm. That is, according to embodiments of the present invention, a compute node may execute a collective operation algorithm that requires a particular number of participants with an arbitrary number of compute nodes and a calculated number of ghost nodes.

For further explanation, FIG. 7 sets forth a flow chart illustrating another example method executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The method of FIG. 7 is similar to the method of FIG. 6 in that the method of FIG. 7 also includes determining (602) a required number (650) of participants for execution of a collective operation algorithm (228); determining (604) a number (652) of contributing nodes (196) having data to participate in the algorithm (228); calculating (606) a number (656) of ghost nodes to participate in the algorithm (228); selecting (608) from a plurality of ghost nodes (197), the calculated number (654) of ghost nodes (198, 199) for participation in the execution of the algorithm (228); and executing (610) the algorithm (228) with both the selected ghost nodes (198, 199) and the contributing nodes (196).

In the method of FIG. 7 also includes mapping (702) a selected ghost node (199) associated with the compute node (297) to physical hardware for execution of the identity operation. Mapping (702) a selected ghost node (199) associated with the compute node (297) to physical hardware for execution of the identity operation may be carried out by assigning a ghost node to one or more resources of a compute node. For further explanation, FIG. 8 sets forth a flow chart illustrating another example method executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The method of FIG. 8 is similar to the method of FIG. 6 in that the method of FIG. 8 also includes determining (602) a required number (650) of participants for execution of a collective operation algorithm (228); determining (604) a number (652) of contributing nodes (196) having data to participate in the algorithm (228); calculating (606) a number (656) of ghost nodes to participate in the algorithm (228); selecting (608) from a plurality of ghost nodes (197), the calculated number (654) of ghost nodes (198, 199) for participation in the execution of the algorithm (228); and executing (610) the algorithm (228) with both the selected ghost nodes (198, 199) and the contributing nodes (196).

In the method of FIG. 8 also includes selecting (802) from a plurality (850) of algorithms, the collective operation algorithm (228) based on the calculated number (654) of ghost nodes to participate in the algorithm (228). Compute nodes may have any number of algorithms to choose from to achieve a particular result. As explained above, the present invention enables a compute node to execute an algorithm that requires a particular number of participants using an arbitrary number of nodes. Selecting (802) from a plurality (850) of algorithms, the collective operation algorithm (228) based on the calcu-

14

lated number (654) of ghost nodes to participate in the algorithm (228) may be carried out by examining the calculated number of ghost nodes needed to 'fill in' during execution of a particular algorithm with an identity operation. For example, a compute node may determine it is more efficient to use a first algorithm that requires two ghost nodes than it is to use a second algorithm that requires a hundred thousand ghost nodes to meet the second algorithm's required number of participants. That is, an algorithm is selected at least partly based on the number of ghost nodes needed to perform the algorithm, including a determination that no ghost nodes are required. Selecting (802) from a plurality (850) of algorithms, the collective operation algorithm (228) based on the calculated number (654) of ghost nodes to participate in the algorithm (228) may also be carried out by sending one or more active messages between the compute nodes of the operational group to determine which algorithm to use.

For further explanation, FIG. 9 sets forth a flow chart illustrating another example method executing a collective operation algorithm in a parallel computer according to embodiments of the present invention. The method of FIG. 9 is similar to the method of FIG. 6 in that the method of FIG. 9 also includes determining (602) a required number (650) of participants for execution of a collective operation algorithm (228); determining (604) a number (652) of contributing nodes (196) having data to participate in the algorithm (228); calculating (606) a number (656) of ghost nodes to participate in the algorithm (228); selecting (608) from a plurality of ghost nodes (197), the calculated number (654) of ghost nodes (198, 199) for participation in the execution of the algorithm (228); and executing (610) the algorithm (228) with both the selected ghost nodes (198, 199) and the contributing nodes (196).

In the method of FIG. 9, however, selecting (608) from a plurality of ghost nodes (197), the calculated number (654) of ghost nodes (198, 199) for participation in the execution of the algorithm (228) includes determining (902), by the compute node (297), to involve a ghost node (199) corresponding to the compute node (297) in participation of execution of the algorithm (228). For example, after selecting the algorithm and determining the number of ghost nodes need for execution of the algorithm, the ghost nodes are selected and invoked. Determining (902), by the compute node (297), to involve a ghost node (199) corresponding to the compute node (297) in participation of execution of the algorithm (228) may be carried out by examining the properties of the algorithm and the operational group. For example, the compute nodes may determine which compute nodes should involve a ghost node based on the 'rank' of the compute node.

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable transmission medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semi-

15

conductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

A computer readable transmission medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable transmission medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of

16

manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

It will be understood from the foregoing description that modifications and changes may be made in various embodiments of the present invention without departing from its true spirit. The descriptions in this specification are for purposes of illustration only and are not to be construed in a limiting sense. The scope of the present invention is limited only by the language of the following claims.

What is claimed is:

1. A method of executing a collective operation algorithm in a parallel computer, the parallel computer comprising a plurality of compute nodes, the compute nodes coupled for data communications by one or more data communications networks, the compute nodes organized in an operational group for collective operations, the method comprising:

determining, by a compute node of an operational group, a required number of participants for execution of a collective operation algorithm;

determining, by the compute node, a number of contributing nodes having data to participate in the algorithm;

calculating, by the compute node, a number of ghost nodes to participate in the algorithm, wherein the number of ghost nodes being the required number of participants minus the number of contributing nodes having data to participate;

selecting, by the compute node, the collective operation algorithm based on the calculated number of ghost nodes;

selecting from a plurality of ghost nodes, by the compute node, the calculated number of ghost nodes for participation in the execution of the algorithm, wherein selecting the calculated number of ghost nodes comprises:

determining, by the compute node, based on properties of the selected collective operation algorithm and rank of contributing nodes, to involve a subset of

17

ghost nodes from the plurality of ghost nodes corresponding to one or more of the contributing nodes; and

executing the algorithm with both the selected ghost nodes and the one or more contributing nodes. 5

2. The method of claim 1 wherein each selected ghost node is configured to participate in the algorithm with an identity operation; and wherein during execution of the algorithm, each selected ghost node participates with the identity operation. 10

3. The method of claim 1 further comprising mapping, by the compute node, a selected ghost node associated with the compute node to physical hardware for execution of an identity operation.

4. The method of claim 1 wherein each compute node of the distributed processing system creates during startup of the compute node, a corresponding ghost node. 15

* * * * *

18